

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188		
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>					
1. REPORT DATE (DD-MM-YYYY) 30-03-2013		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 15-Sep-2008 - 14-Sep-2012	
4. TITLE AND SUBTITLE Final Report: Analysis, Evaluation and Improvement of Sequential Single-Item Auctions for the Cooperative Real-Time Allocation of Tasks			5a. CONTRACT NUMBER W911NF-08-1-0468		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER 611102		
6. AUTHORS Sven Koenig, Craig Tovey			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES University of Southern California Department of Contracts and Grants 3720 S. Flower St. Los Angeles, CA 90089 -0701			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 55072-NS.51		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT Auctions and similar market-based mechanisms are decentralized methods that are a promising way for teams of agents to assign and re-assign tasks among themselves in dynamic, partially known and time-constrained domains. We report on the results of our project on the "Analysis, Evaluation and Improvement of Sequential Single-Item Auctions for the Cooperative Real-Time Allocation of Tasks," where we designed efficient and effective auction mechanisms and analyzed their properties.					
15. SUBJECT TERMS Agent Coordination, Cooperative Auctions, Market-Based Mechanisms, Pursuit Evasion					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Sven Koenig
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			19b. TELEPHONE NUMBER 213-740-6491

Report Title

Final Report: Analysis, Evaluation and Improvement of Sequential Single-Item Auctions for the Cooperative Real-Time Allocation of Tasks

ABSTRACT

Auctions and similar market-based mechanisms are decentralized methods that are a promising way for teams of agents to assign and re-assign tasks among themselves in dynamic, partially known and time-constrained domains. We report on the results of our project on the "Analysis, Evaluation and Improvement of Sequential Single-Item Auctions for the Cooperative Real-Time Allocation of Tasks," where we designed efficient and effective auction mechanisms and analyzed their properties.

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

<u>Received</u>	<u>Paper</u>
03/30/2013	19.00 Richard Borie, Craig Tovey, Sven Koenig. Algorithms and Complexity Results for Graph-Based Pursuit Evasion, Autonomous Robots , (12 2011): 317. doi:
03/30/2013	29.00 Xiaoming Zheng, Sven Koenig, David Kempe, Sonal Jain. Multi-Robot Forest Coverage for Weighted and Unweighted Terrain, IEEE Transactions on Robotics, (12 2010): 1018. doi:
03/30/2013	30.00 Kenneth Daniel, Alex Nash, Sven Koenig, Ariel Felner. Theta*: Any-Angle Path Planning on Grids, Journal on Artificial Intelligence Research, (12 2010): 533. doi:
03/30/2013	31.00 Craig Tovey, Sven Koenig. Localization: Approximation and Performance Bounds to Minimize Travel Distance, IEEE Transactions on Robotics, (12 2010): 320. doi:
03/30/2013	32.00 William Yeoh, Ariel Felner, Sven Koenig. BnB-ADOPT: An Asynchronous Branch-and-Bound DCOP Algorithm, Journal on Artificial Intelligence Research, (12 2010): 85. doi:
TOTAL:	5

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

<u>Received</u>	<u>Paper</u>
-----------------	--------------

03/30/2013	18.00	Sven Koenig. Making Good Decisions Quickly, IEEE Intelligent Informatics Bulletin , (12 2012): 14. doi:
------------	-------	--

TOTAL:	1
---------------	----------

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Number of Presentations:	0.00
---------------------------------	------

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

<u>Received</u>	<u>Paper</u>
-----------------	--------------

03/30/2013	23.00	Tansel Uras, Sven Koenig, Carlos Hernandez. Subgoal Graphs for Eight-Neighbor Gridworlds, Symposium on Combinatorial Search. , . : ,
------------	-------	---

03/30/2013	35.00	Sven Koenig. Creating a Uniform Framework for Task and Motion Planning: A Case for Incremental Heuristic Search? [Overview Paper], ICAPS-10 Workshop on Combining Action and Motion Planning. , . : ,
------------	-------	---

TOTAL:	2
---------------	----------

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

<u>Received</u>	<u>Paper</u>
03/30/2013	20.00 Xiaoming Zheng, Sven Koenig. Generalized Reaction Functions for Solving Complex-Task Allocation Problems, International Joint Conference on Artificial Intelligence (IJCAI). , . : ,
03/30/2013	21.00 Carlos Hernandez, Jorge Baier, Tansel Uras, Sven Koenig. Time-Bounded Adaptive A*, International Joint Conference on Autonomous Agents and Multiagent Systems. , . : ,
03/30/2013	22.00 Carlos Hernandez, Jorge Baier, Tansel Uras, Sven Koenig. Position Paper: Incremental Search Algorithms Considered Poorly Understood, Position Paper: Incremental Search Algorithms Considered Poorly Understood. , . : ,
03/30/2013	24.00 Xiaoxun Sun, Tansel Uras, Sven Koenig, William Yeoh. Incremental ARA*: An Incremental Anytime Search Algorithm for Moving-Target Search, International Conference on Automated Planning and Scheduling (ICAPS). , . : ,
03/30/2013	26.00 Carlos Hernandez, Xiaoxun Sun, Sven Koenig, Pedro Meseguer. Tree Adaptive A* , International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS). , . : ,
03/30/2013	27.00 William Yeoh, Pradeep Varakantham, Xiaoxun Sun, Sven Koenig. Incremental DCOP Search Algorithms for Solving Dynamic DCOPs [Extended Abstract], Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS). , . : ,
03/30/2013	33.00 Alex Nash, Sven Koenig, Craig Tovey. Lazy Theta*: Any-Angle Path Planning and Path Length Analysis in 3D, AAAI Conference on Artificial Intelligence (AAAI). , . : ,
03/30/2013	34.00 Xiaoming Zheng, Sven Koenig. Market-Based Algorithms for Allocating Complex Tasks [Student Abstract], Proceedings of the AAAI Conference on Artificial Intelligence (AAAI). , . : ,
03/30/2013	36.00 Yilliam Yeoh, Roie Zivan, Sven Koenig. Discrepancy-Based Approach for Solving Distributed Constraint Optimization Problems, International Workshop on Distributed Constraint Reasoning (DCR). , . : ,
03/30/2013	37.00 Xiaoxun Sun, William Yeoh, Sven Koenig. Generalized Fringe-Retrieving A*: Faster Moving Target Search on State Lattices, International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS). , . : ,
03/30/2013	38.00 Xiaoxun Sun, William Yeoh, Sven Koenig. Moving Target D* Lite, International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS). , . : ,
03/30/2013	39.00 Kenneth Daniel, Richard Borie, Sven Koenig, Craig Tovey. ESP: Pursuit Evasion on Series-Parallel Graphs [Poster Abstract], International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS). , . : ,
03/30/2013	40.00 Xiaoming Zheng, Sven Koenig. Negotiation with Reaction Functions for Solving Complex Task Allocation Problems, International Conference on Intelligent Robots and Systems (IROS). , . : ,
03/30/2013	41.00 Richard Borie, Craig Tovey, Sven Koenig. Algorithms and Complexity Results for Pursuit-Evasion Problems, International Joint Conference on Artificial Intelligence (IJCAI). , . : ,

03/30/2013	42.00	Alex Nash, Sven Koenig, Maxim Likhachev. Incremental Phi*: Incremental Any-Angle Path Planning on Grids, International Joint Conference on Artificial Intelligence (IJCAI). , . : ,
03/30/2013	43.00	Xiaoming Zheng, Sven Koenig. K-Swaps: Cooperative Negotiation for Solving Task-Allocation Problems, International Joint Conference on Artificial Intelligence (IJCAI). , . : ,
03/30/2013	44.00	Xiaoxun Sun, William Yeoh, Sven Koenig. Efficient Incremental Search for Moving Target Search, International Joint Conference on Artificial Intelligence (IJCAI). , . : ,
03/30/2013	45.00	William Yeoh, Xiaoxun Sun, Sven Koenig. Trading Off Solution Quality for Faster Computation in DCOP Search Algorithms, International Joint Conference on Artificial Intelligence (IJCAI). , . : ,
03/30/2013	46.00	Ali Ekici, Pinar Keskinocak, Sven Koenig. Multi-Robot Routing with Linear Decreasing Rewards over Time, IEEE International Conference on Robotics and Automation (ICRA). , . : ,
03/30/2013	47.00	Carlos Hernandez, Pedro Meseguer, Xiaoxun Sun, Sven Koenig. Path-Adaptive A* for Incremental Heuristic Search in Unknown Terrain, International Conference on Automated Planning and Scheduling (ICAPS). , . : ,
03/30/2013	48.00	Sven Koenig. Open Problem: Analyzing Ant Robot Coverage, International Conference on Learning Theory (COLT). , . : ,
03/30/2013	49.00	Sven Koenig, Pinar Keskinocak , Craig Tovey. Progress on Agent Coordination with Cooperative Auctions [Senior Member Paper], Proceedings of the AAAI Conference on Artificial Intelligence (AAAI). , . : ,
03/30/2013	50.00	Xiaoming Zheng, Sven Koenig. Sequential Incremental-Value Auctions, AAAI Conference on Artificial Intelligence (AAAI). , . : ,

TOTAL: 23

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts):

(d) Manuscripts

<u>Received</u>	<u>Paper</u>
-----------------	--------------

03/05/2010	9.00	K. Daniel, R. Borie, S. Koenig, C. Tovey. ESP: Pursuit Evasion on Series-Parallel Graphs [Poster Abstract], Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2010 (03 2010)
03/05/2010	11.00	X. Sun, W. Yeoh, S. Koenig. Generalized Fringe-Retrieving A*: Faster Moving Target Search on State Lattices, (03 2010)
03/05/2010	10.00	X. Sun, W. Yeoh, S. Koenig. Moving Target D* Lite*, Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010) (03 2010)
03/23/2009	1.00	S. Koenig. Multi-Robot Routing with Linear Decreasing Rewards over Time, ()
04/24/2009	3.00	X. Sun, W. Yeoh, S. Koenig. Efficient Incremental Search for Moving Target Search, ()
04/24/2009	2.00	W. Yeoh, X. Sun, S. Koenig. Trading Off Solution Quality for Faster Computation in DCOP Search Algorithms, ()
04/24/2009	4.00	X. Zheng, S. Koenig. K-Swaps: Cooperative Negotiation for Solving Task-Allocation Problems, ()
04/24/2009	5.00	A. Nash, S. Koenig, M. Likhachev. Incremental Phi: Incremental Any-Angle Path Planning on Grids, ()
04/24/2009	6.00	R. Borie, C. Tovey, S. Koenig. Algorithms and Complexity Results for Pursuit-Evasion Problems, ()
08/19/2009	8.00	Xiaoming Zheng, Sven Koenig. Negotiation with Reaction Functions for Solving Complex Task Allocation Problems, Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (2009)

TOTAL: 10

Number of Manuscripts:

Books

<u>Received</u>	<u>Paper</u>
-----------------	--------------

TOTAL:

Patents Submitted

Patents Awarded

Awards

Some of our general honors (related to all of our research, not necessarily specifically to research from this particular grant) from 2009 to today include:

- The entry “Eight-Neighbor Gridworlds” by T. Uras, S. Koenig and C. Hernandez was declared a non-dominated optimal entry by the Grid-Based Path Planning Competition in 2012 (note: preliminary result).
- Sven Koenig gave keynote speeches at the IEEE/WIC/ACM International Conference on Intelligent Agent Technology in 2012 and at the German Conference on Artificial Intelligence in 2011. He also received a Certificate of Appreciation from the Web Intelligence Consortium in 2012.
- Sven Koenig was promoted to full professor in 2011.
- Sven Koenig gave an invited talk in the Distinguished Lecturer Seminar of the Computer Science Department of UC Irvine in 2010 on Market Mechanisms for the Allocation of Resources in Cooperative Domains.
- Sven Koenig received the Mellon Award for Faculty Mentoring Undergraduate Students in 2009.
- Sven Koenig gave several invited talks at conferences, including the ICAPS Workshop on Combining Task and Motion Planning for Real-World Applications in 2012, the AAAI Fall Symposium Series: Workshop on Multi-Agent Coordination under Uncertainty in 2011, the AAMAS Workshop on Optimisation in Multi-Agent Systems in 2011, the Annual Consortium for Computing Sciences in Colleges - Southwestern Regional Conference in 2011, the AAAI Workshop on Bridging the Gap between Task and Motion Planning in 2010, the International Symposium on Combinatorial Search in 2010, the Dagstuhl Seminar on Cognitive Robotics in 2010, and the ICRA Workshop on Search and Pursuit/Evasion in the Physical World: Efficiency, Scalability, and Guarantees in 2010.
- Sven Koenig became a member of the advisory board of the National ICT Australia Optimization Group in 2011, a member of the governing council of the International Symposium on Combinatorial Search in 2010 and an associate editor of the Autonomous Agents and Multi-Agent Systems journal in 2010. He was also co-chair of the International Symposium on Combinatorial Search in 2009.
- Craig Tovey received the Class of 1934 Outstanding Interdisciplinary Activities Award from Georgia Institute of Technology in 2011.
- David Vandegrift and Craig Tovey received a 2010 finalist award in the SAIC-Georgia Institute of Technology Student Paper Competition.

Graduate Students

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	Discipline
Xiaoxun Sun	0.00	
Kenneth Daniel	0.00	
Tansel Uras	0.00	
Alex Nash	0.00	
Yan Shu	0.00	
FTE Equivalent:	0.00	
Total Number:	5	

Names of Post Doctorates

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
Masabumi Furuhashi	0.00
FTE Equivalent:	0.00
Total Number:	1

Names of Faculty Supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	National Academy Member
Sven Koenig	0.00	
Craig Tovey	0.00	
FTE Equivalent:	0.00	
Total Number:	2	

Names of Under Graduate students supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	Discipline
Daniel Vandegrift	0.00	Industrial and Systems Engineering
FTE Equivalent:	0.00	
Total Number:	1	

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period: 1.00

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:..... 1.00

The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:..... 0.00

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):..... 1.00

Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense 0.00

The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields: 0.00

Names of Personnel receiving masters degrees

<u>NAME</u>
Kenneth Daniel
Total Number:

1

Names of personnel receiving PHDs

<u>NAME</u>
Xiaoxun Sun (pending)
Alex Nash
Yan Shu
Total Number:

3

Names of other research staff

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

Technology Transfer

Final Report

Auctions and similar market-based mechanisms are decentralized methods that are a promising way for teams of agents to assign and re-assign tasks among themselves in dynamic, partially known and time-constrained domains. In this case, the agents are the bidders and the tasks are the goods up for auction. Every agent then executes the tasks that it wins. As the agents discover more about the environment during execution or as the environment or tasks change, they run additional auctions to re-optimize the allocation of tasks among themselves. Auction-based coordination systems promise to be efficient in terms of communication (agents communicate only essential information, such as numeric bids which encapsulate large amounts of local information) and computation (agents compute their bids in parallel).

Our project was an interdisciplinary effort of two researchers who have already worked with each other productively on other topics. One of the PIs was a researcher in artificial intelligence from the University of Southern California, and the other PI was a researcher in theoretical computer science and operations research from Georgia Institute of Technology, which enables us to build systems with very good performance on top of sound mathematical foundations. In particular, we designed efficient and effective auction mechanisms for the cooperative real-time allocation of tasks and applied the results to simulations of mine clearing tasks and search-and-rescue tasks. In general, auction-based coordination systems based on sequential single-item auctions result in no more bids and provide much better performance guarantees than auction-based coordination systems based on parallel single-item auctions. They are also much easier to implement than auction-based coordination systems based on combinatorial auctions (since the number of bids is much smaller and winner determination is much simpler and faster) and allow for a decentralized implementation. We therefore designed efficient and effective auction mechanisms for the cooperative real-time allocation of tasks based on sequential single-item auctions and improved them in several ways, bringing to bear insights from different disciplines, namely robotics and artificial intelligence on one hand and theoretical computer science and operations research on the other hand. We also analyzed their properties theoretically and evaluated them experimentally to characterize when they should be used and how good they are, creating a mathematical framework for the design and analysis of auctions for the cooperative real-time allocation of tasks.

While there is interest in the research community in using auctions for the coordination of teams of agents, most of the work had been empirical and somewhat ad-hoc in nature. Auctions have, of course, long been studied in economics. However, the auction literature in economics assumes a competitive context with rational, and hence long, decision cycles rather than a cooperative context with extreme time pressure. Thus, while some of its insights could be exploited for our purposes (for example, the concept of synergy), most of them did not apply (for example, incentive compatibility to elicit sincere bids) since one controls both the auction mechanism and the preference structures of the agents.

It is often crucial to coordinate teams of cooperating agents well. Auction-based coordination applies to a wide range of real-time domains with multiple agents and tasks, where there are positive and negative synergies among tasks, including on-line distributed routing, role assignment, scheduling or control. The simplest scenario that we studied was on-line distributed routing, where a set of targets is given and each one of them needs to be visited by one robot. On-line distributed routing includes assigning accidents to ambulances, incidents to police cars, customers to taxi cabs, rocks to Mars rovers (to take

and analyze rock probes), observation locations to robots that map terrain or wings of an art gallery to guard robots. Thus, it applies directly to tasks such as mine clearing and search-and-rescue, which we used as both motivating scenarios and testbeds. Our experimental results show that it is feasible to allocate hundreds of tasks in real-time. We also performed experiments in an existing search-and-rescue simulation to be relevant to applications.

Overall, our research showed that auctions appear to be well suited for the cooperative real-time allocation of tasks and can be used as foundation for building coordination systems. There now exists the beginning of a theoretical framework for understanding their properties and for extending their capabilities to more realistic settings. Our findings are summarized in our publications, listed below. We provide the abstracts of some key publications to give the reader an overview of our main results and insights. We have also attached a copy of our paper

S. Koenig, P. Keskinocak and C. Tovey. Progress on Agent Coordination with Cooperative Auctions [Senior Member Paper]. In Proceedings of the AAAI Conference on Artificial Intelligence (AAAI), 2010.

which gives a good overview of our main results and insights up to and including 2010.

We also studied pursuit-evasion problems where a number of pursuers have to clear a given graph of potentially fast-moving evaders despite poor visibility, for example, where robots search a cave system to find criminals. We studied when polynomial-time algorithms exist to determine how many robots are needed to clear a given graph and how a given number of robots should move on the graph to clear it with either a minimum sum of their travel distances or minimum task-completion time.

Key Publications

In the following, we give the abstracts of some of our key publications. These publications acknowledge ARO funding as well as funding from other grants. The full text of most of these publications is available at idm-lab.org/publications.html.

2012

S. Koenig. Making Good Decisions Quickly. The IEEE Intelligent Informatics Bulletin, 13, (1), 14-20, 2012.

Abstract: Several disciplines, including artificial intelligence, operations research and many others, study how to make good decisions. In this overview article, we argue that the key to making progress in our research area is to combine their ideas, which often requires serious technical advances to reconcile their different assumptions and methods in a way that results in synergy among them. To illustrate this point, we give a broad overview of our ongoing research on search and planning (with a large number of students and colleagues, both at the University of Southern California and elsewhere) to demonstrate how to combine ideas from different decision making disciplines. For example, we describe how to combine ideas from artificial intelligence, operations research, and utility theory to create the foundations for building decision support systems that fit the risk preferences of human decision makers in high-stake one-shot decision situations better than current systems. We also describe how to combine ideas from artificial intelligence, economics, theoretical computer science and operations

research to build teams of robots that use auctions to distribute tasks autonomously among themselves, and give several more examples.

2011

R. Borie, C. Tovey and S. Koenig. Algorithms and Complexity Results for Graph-Based Pursuit Evasion. *Autonomous Robots*, 31, (4), 317-332, 2011.

Abstract: We study the classical edge-searching pursuit-evasion problem where a number of pursuers have to clear a given graph of fast-moving evaders despite poor visibility, for example, where robots search a cave system to ensure that no terrorists are hiding in it. We study when polynomial-time algorithms exist to determine how many robots are needed to clear a given graph (minimum robot problem) and how a given number of robots should move on the graph to clear it with either a minimum sum of their travel distances (minimum distance problem) or minimum task-completion time (minimum time problem). The robots cannot clear a graph if the vertex connectivity of some part of the graph exceeds the number of robots. Researchers therefore focus on graphs whose subgraphs can always be cut at a limited number of vertices, that is, graphs of low treewidth, typically trees. We describe an optimal polynomial-time algorithm, called CLEARHETREE, that is shorter and algorithmically simpler than the state-of-the-art algorithm for the minimum robot problem on unit-width unit-length trees. We then generalize prior research to both unit-width arbitrary-length and unit-length arbitrary-width graphs and derive both algorithms and time complexity results for a variety of graph topologies. Pursuit-evasion problems on the former graphs are generally simpler than pursuit-evasion problems on the latter graphs. For example, the minimum robot and distance problems are solvable in polynomial time on unit-width arbitrary-length trees but NP-hard on unit-length arbitrary-width trees.

X. Zheng and S. Koenig. Generalized Reaction Functions for Solving Complex-Task Allocation Problems. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, pages 478-483, 2011.

Abstract: We study distributed task-allocation problems where cooperative agents need to perform some tasks simultaneously. Examples are multi-agent routing problems where several agents need to visit some targets simultaneously, for example, to move obstacles out of the way cooperatively. In this paper, we first generalize the concept of reaction functions proposed in the literature to characterize the agent costs of performing multiple complex tasks. Second, we show how agents can construct and approximate reaction functions in a distributed way. Third, we show how reaction functions can be used by an auction-like algorithm to allocate tasks to agents. Finally, we show empirically that the team costs of our algorithms are substantially smaller than those of an existing state-of-the-art allocation algorithm for complex tasks.

2010

X. Zheng and S. Koenig. Sequential Incremental-Value Auctions. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2010.

Abstract: We study the distributed allocation of tasks to cooperating robots in real time, where each task has to be assigned to exactly one robot so that the sum of the latencies of all tasks is as small as possible. We propose a new auction-like algorithm, called Sequential Incremental-Value (SIV) auction, which assigns tasks to robots in multiple rounds. The idea behind SIV auctions is to assign as many tasks

per round to robots as possible as long as their individual costs for performing these tasks are at most a given bound, which increases exponentially from round to round. Our theoretical results show that the team costs of SIV auctions are at most a constant factor larger than minimal.

S. Koenig, P. Keskinocak and C. Tovey. Progress on Agent Coordination with Cooperative Auctions [Senior Member Paper]. In Proceedings of the AAAI Conference on Artificial Intelligence (AAAI), 2010.

Abstract: Auctions are promising decentralized methods for teams of agents to allocate and re-allocate tasks among themselves in dynamic, partially known and time-constrained domains with positive or negative synergies among tasks. Auction-based coordination systems are easy to understand, simple to implement and broadly applicable. They promise to be efficient both in communication (since agents communicate only essential summary information) and in computation (since agents compute their bids in parallel). Artificial intelligence research has explored auction-based coordination systems since the early work on contract networks [Smith1980], mostly from an experimental perspective. This overview paper describes our recent progress towards creating a framework for the design and analysis of cooperative auctions for agent coordination.

K. Daniel, R. Borie, S. Koenig and C. Tovey. ESP: Pursuit Evasion on Series-Parallel Graphs [Poster Abstract]. In Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2010.

Abstract: We develop a heuristic approach, called ESP, that solves large pursuit-evasion problems on series-parallel (that is, treewidth-2) graphs quickly and with small costs. It exploits their topology by performing dynamic programming on their decomposition graphs. We show that ESP scales up to much larger graphs than a strawman approach based on previous results from the literature.

2009

R. Borie, C. Tovey and S. Koenig. Algorithms and Complexity Results for Pursuit-Evasion Problems. In Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI), pages 59-66, 2009.

Abstract: We study pursuit-evasion problems where a number of pursuers have to clear a given graph. We study when polynomial-time algorithms exist to determine how many pursuers are needed to clear a given graph and how a given number of pursuers should move on the graph to clear it with either a minimum sum of their travel distances or minimum task-completion time. We generalize prior work to both unit-width arbitrary-length and unit-length arbitrary-width graphs and derive both algorithms and complexity results for a variety of graph topologies. In this context, we describe a polynomial-time algorithm, called CLEARHETREE, that is much shorter and algorithmically simpler than the state-of-the-art algorithm for the minimum pursuer problem on trees. Our theoretical research lays a firm theoretical foundation for pursuit evasion on graphs and informs practitioners about which problems are easy and which ones are hard.

X. Zheng and S. Koenig. K-Swaps: Cooperative Negotiation for Solving Task-Allocation Problems. In Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI), pages 373-379, 2009.

Abstract: In this paper, we study distributed algorithms for cooperative agents that allow them to exchange their assigned tasks in order to reduce their team cost. We define a new type of contract,

called K-swaps, that describes multiple task exchanges among multiple agents at a time, which generalizes the concept of single task exchanges. We design a distributed algorithm that constructs all possible K-swaps that reduce the team cost of a given task allocation and show that each agent typically only needs to communicate a small part of its local computation results to the other agents. We then demonstrate empirically that K-swaps can reduce the team costs of several existing task-allocation algorithms significantly even if K is small.

A. Ekici, P. Keskinocak and S. Koenig. Multi-Robot Routing with Linear Decreasing Rewards over Time. In Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), pages 958-963, 2009.

Abstract: We study multi-robot routing problems (MR-LDR) where a team of robots has to visit a set of given targets with linear decreasing rewards over time, such as required for the delivery of goods to rescue sites after disasters. The objective of MR-LDR is to find an assignment of targets to robots and a path for each robot that maximizes the surplus, which is defined to be the total reward collected by the team minus its total travel cost. We develop a mixed integer program that solves MR-LDR optimally with a flow-type formulation and can be solved faster than the standard TSP-type formulations but also show that solving MR-LDR optimally is NP-hard. We then develop an auction-based algorithm and demonstrate that it solves MR-LDR in seconds and with a surplus that is comparable to the surplus found by the mixed integer program with a 12 hour time limit.

X. Zheng and S. Koenig. Negotiation with Reaction Functions for Solving Complex Task Allocation Problems. In Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS), pages 4811-4816, 2009.

Abstract: We study task-allocation problems where cooperative robots need to perform tasks simultaneously. We develop a distributed negotiation procedure that allows robots to find all task exchanges that reduce the team cost of a given task allocation, without robots having to know how other robots compute their robot costs. Finally, we demonstrate empirically that our negotiation procedure can substantially reduce the team costs of task allocations resulting from existing task-allocation procedures, including sequential single-item auctions.

Additional Publications

Other publications that acknowledge ARO funding (as well as funding from other grants) include the following ones. The full text of most of these publications is available at idm-lab.org/publications.html. Dissertations are also listed but might not acknowledge ARO funding.

2012

C. Hernandez, J. Baier, T. Uras and S. Koenig, Time-Bounded Adaptive A*, Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), pages 997-1006, 2012. A short version of this paper appeared also in: Proceedings of the Symposium on Combinatorial Search (SoCS), 2012. A summary of this paper appeared also in: C. Hernandez, J. Baier, T. Uras and S. Koenig, New Developments in Real-Time Heuristic Search: A Demo [System Demonstration

and Exhibition Abstract], Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS), 2012.

C. Hernandez, J. Baier, T. Uras and S. Koenig. Position Paper: Incremental Search Algorithms Considered Poorly Understood. In Proceedings of the Symposium on Combinatorial Search, 2012.

T. Uras, S. Koenig and C. Hernandez. Subgoal Graphs for Eight-Neighbor Gridworlds. In Proceedings of the Symposium on Combinatorial Search, 2012.

X. Sun, T. Uras, S. Koenig and W. Yeoh. Incremental ARA*: An Incremental Anytime Search Algorithm for Moving-Target Search. In Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS), 2012.

A. Nash. Any-Angle Path Planning. Dissertation. University of Southern California, 2012.

2011

C. Hernandez, X. Sun, S. Koenig and P. Meseguer. Tree Adaptive A*. In Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), pages 123-130, 2011.

W. Yeoh, P. Varakantham, X. Sun and S. Koenig. Incremental DCOP Search Algorithms for Solving Dynamic DCOPs [Extended Abstract]. In Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), pages 169-170, 2011.

Y. Shu, Future Aircraft Networks and schedules. Dissertation. Georgia Institute of Technology, 2011.

2010

X. Zheng, S. Koenig, D. Kempe and S. Jain. Multi-Robot Forest Coverage for Weighted and Unweighted Terrain. IEEE Transactions on Robotics, 26, (6), 1018-1031, 2010.

K. Daniel, A. Nash, S. Koenig and A. Felner. Theta*: Any-Angle Path Planning on Grids. Journal of Artificial Intelligence Research, 39, 533-579, 2010.

C. Tovey and S. Koenig. Localization: Approximation and Performance Bounds to Minimize Travel Distance. IEEE Transactions on Robotics, 26, (2), 320-330, 2010.

W. Yeoh, A. Felner and S. Koenig. BnB-ADOPT: An Asynchronous Branch-and-Bound DCOP Algorithm. Journal of Artificial Intelligence Research, 38, 85-133, 2010.

A. Nash, S. Koenig and C. Tovey. Lazy Theta*: Any-Angle Path Planning and Path Length Analysis in 3D. In Proceedings of the AAAI Conference on Artificial Intelligence (AAAI), 2010.

X. Zheng and S. Koenig. Market-Based Algorithms for Allocating Complex Tasks [Student Abstract]. In Proceedings of the AAAI Conference on Artificial Intelligence (AAAI), 2010.

X. Sun, W. Yeoh and S. Koenig. Generalized Fringe-Retrieving A*: Faster Moving Target Search on State Lattices. In Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), pages 1081-1088, 2010.

X. Sun, W. Yeoh and S. Koenig. Moving Target D* Lite. In Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), pages 67-74, 2010.

S. Koenig. Creating a Uniform Framework for Task and Motion Planning: A Case for Incremental Heuristic Search? [Overview Paper]. In Proceedings of the ICAPS-10 Workshop on Combining Action and Motion Planning, 2010.

S. Koenig. Open Problem: Analyzing Ant Robot Coverage. In Proceedings of the International Conference on Learning Theory (COLT), pages 312-313, 2010.

2009

A. Nash, S. Koenig and M. Likhachev, Incremental Phi*: Incremental Any-Angle Path Planning on Grids, Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI), page 1824-1830, 2009. A version of this paper was also presented in: International Symposium on Combinatorial Search (SoCS), 2009.

X. Sun, W. Yeoh and S. Koenig, Efficient Incremental Search for Moving Target Search, Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI), pages 615-620, 2009. A short version of this paper appeared also in: Proceedings of the ICAPS-09 Doctoral Consortium, pages 29-32, 2009.

W. Yeoh, X. Sun and S. Koenig, Trading Off Solution Quality for Faster Computation in DCOP Search Algorithms, Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI), pages 354-360, 2009. A version of this paper was also presented in: International Symposium on Combinatorial Search (SoCS), 2009. Earlier versions of this paper appeared also in: W. Yeoh, S. Koenig and X. Sun, Trading Off Solution Cost for Smaller Runtime in DCOP Search Algorithms [Poster Abstract], Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), pages 1445-1448, 2008. Proceedings of the Tenth International Workshop on Distributed Constraint Reasoning (DCR), pages 25-35, 2008.

W. Yeoh, R. Zivan and S. Koenig. Discrepancy-Based Approach for Solving Distributed Constraint Optimization Problems. In Proceedings of the International Workshop on Distributed Constraint Reasoning (DCR), pages 132-144, 2009.

C. Hernandez, P. Meseguer, X. Sun and S. Koenig. Path-Adaptive A* for Incremental Heuristic Search in Unknown Terrain. In Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS), pages 358-361, 2009.

Progress on Agent Coordination with Cooperative Auctions*

Sven Koenig

University of Southern California
Computer Science
skoening@usc.edu

Pinar Keskinocak and Craig Tovey

Georgia Institute of Technology
Industrial and Systems Engineering
{pinar,ctovey}@isye.gatech.edu

Abstract

Auctions are promising decentralized methods for teams of agents to allocate and re-allocate tasks among themselves in dynamic, partially known and time-constrained domains with positive or negative synergies among tasks. Auction-based coordination systems are easy to understand, simple to implement and broadly applicable. They promise to be efficient both in communication (since agents communicate only essential summary information) and in computation (since agents compute their bids in parallel). Artificial intelligence research has explored auction-based coordination systems since the early work on contract networks (Smith 1980), mostly from an experimental perspective. This overview paper describes our recent progress towards creating a framework for the design and analysis of cooperative auctions for agent coordination.

Introduction

Centralized control is often inefficient for distributed systems in terms of both the required amount of computation and communication since the central controller is the bottleneck of the system. Many researchers have therefore studied agent coordination with cooperative auctions. An auction is “a market institution with an explicit set of rules determining resource allocation and prices on the basis of bids from the market participants” (McAfee and McMillan 1987). In auction-based coordination systems, the bidders are agents, and the items up for auction are tasks to be executed by the agents. All agents bid their costs. Thus, the agent with the smallest bid cost is best suited for a task. All agents then execute the tasks that they win. Economics has an extensive auction literature but its agents are rational and com-

petitive, leading to long decision cycles, strategic behavior and possibly collusion. Such issues do not arise in auction-based coordination systems because the agents simply execute their program. Thus, while some insights from economics can be exploited for building auction-based coordination systems (for example, the concept of synergy), many of them do not apply. Conversely, auction-based coordination systems must operate in real-time. Some researchers therefore prefer to use the term “auction-inspired control algorithms” (for decentralized control) over “cooperative auctions” to highlight these differences. This paper provides a unified overview of our progress towards creating a framework for the design and analysis of cooperative auctions for agent coordination, drawing from publications in different venues, including robotics and agents conferences.

Applications

Auction-based coordination systems apply to a wide range of real-time domains.

On-Line Distributed Role Allocation

Allocate roles to agents with different capabilities for the execution of a given plan or playbook so that every task is performed by a qualified agent (Hunsberger and Grosz 2000). Examples include allocating attacker and defender roles to robots in RoboSoccer (Frias-Martinez, Sklar, and Parsons 2004), tasks to ambulance teams, fire brigades and police forces in RoboCup Rescue (Nair et al. 2002), tasks to team members that prevent incursions of aircraft or boats into launch impact zones (Sycara et al. 2005), different observation targets to sensors in a wireless sensor network (for example, for bird habitat sensing) (Yu and Prasanna 2005; Howard and Viguria 2007), and observer and manipulator roles for manipulation tasks (such as in box pushing) (Gerkey and Matarić 2002).

On-Line Distributed Scheduling and Control

Allocate tasks (or processes) to machines (or processors) in a distributed system to minimize latency, maximize throughput or optimize other objectives, possibly considering constraints among the tasks, such as precedence constraints or assignment restrictions. Examples include allocating complex workflows with deadlines to CPUs in grid computing

*We thank Xiaoming Zheng for his overview of auction applications and contributions to the development of algorithms and their experimental evaluation. This material is based on some of our published research on the topic, which was supported by, or in part by, NSF under contract/grant number 0413196, ARL/ARO under contract/grant number W911NF-08-1-0468 and ONR in the form of a MURI under contract/grant number N00014-09-1-1031. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the sponsoring organizations, agencies or the U.S. government.
Copyright © 2010, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

(Reeves et al. 2005) and on-line control of environmental conditions in smart buildings (Clearwater et al. 1996).

On-Line Distributed Routing

Allocate locations to agents. Examples include allocating accidents to ambulances, incidents to police cars, mines to autonomous underwater vehicles (for de-mining) (Sariel, Balch, and Stack 2006), search-and-rescue locations to first responders, customers to taxi cabs, rocks to Mars rovers (for taking and analyzing rock probes) (Tovey et al. 2005), mines to submarines (for identification) (Sariel, Balch, and Stack 2006), observation locations to robots that map terrain (Simmons et al. 2000) and wings of an art gallery to guard robots (Kalra, Stentz, and Ferguson 2005). In this context, auctions have been used on actual robots (Thayer et al. 2000; Gerkey and Matarić 2002; Zlot et al. 2002).

Testbed: Multi-Robot Routing

The standard testbed of auction-based coordination systems is multi-robot routing (Dias et al. 2006), a special case of on-line distributed routing. For multi-robot routing, the bidders are the robots, and the items up for auction are tasks to visit given targets (locations). The robots are identical and know both their own location and the target locations. They can move and broadcast information without error. They might initially not know where the obstacles are in the terrain but always observe the ones in their vicinity without error. For ease of exposition, we assume in this overview that the robots have to visit all targets (and do not need to return to their initial locations) with a small sum of travel distances. Multi-robot routing problems are similar to traveling salesperson problems and their variants (Lawler et al. 1985), which simplifies their analysis.

Auction-Based Coordination Systems

Minimizing the sum of travel distances is NP-hard for multi-robot routing problems even if the terrain and targets are initially known and do not change (Lagoudakis et al. 2005). Auction-based coordination systems work as follows: Every robot determines a path with the smallest travel distance to visit all of the (unvisited) targets that are allocated to it and starts to move along the path. Thus, the robot does not necessarily visit the targets in the order in which they were allocated to it. Whenever a robot gains more information about the terrain or observes that the terrain has changed, it shares this information with the other robots. If that information increases the travel distance of at least one robot or new targets are introduced, then all robots put their (unvisited) targets up for auction for re-allocation. Each robot then bids in light of the new information, assuming that no changes will be necessary in the future. The auction closes after a predetermined amount of time and the robots are allocated new targets. Auction-based coordination systems based on *parallel auctions* allocate the targets in independent and simultaneous single-round auctions, one for each target. Every robot bids on all targets. They do not take synergies among targets into account, which often results in

a large sum of travel distances.¹ Auction-based coordination systems based on *combinatorial auctions* allocate the targets in one single-round auction. Every robot bids on all bundles (sets) of targets. Hence, all synergies among targets are taken into account, which minimizes the sum of travel distances. However, an exponential (in the number of targets) number of bids must be generated, transmitted, and processed. The approximations necessary to guarantee real-time performance (Berhault et al. 2003) often interact in unpredictable ways and result in complicated code that is difficult to debug, maintain and integrate into robot architectures. Sequential single-item (SSI) auction-based coordination systems have recently emerged as a promising way of combining the advantages of auction-based coordination systems based on parallel auctions (namely, their small number of bids and fast winner determination) and combinatorial auctions (namely, their small sum of travel distances) (Lagoudakis et al. 2005).

Sequential Single-Item (SSI) Auctions

The targets are allocated in one multi-round auction. During each round, every robot bids on each unallocated target, and winner determination then allocates one additional target to one robot. Every robot bids on a target the smallest increase in its travel distance that would result from it being allocated the target that it bids on in addition to all targets allocated to it in previous rounds (marginal-cost bidding) (Sandholm 1993). Winner determination determines the bid with the smallest bid cost and allocates the corresponding target to the corresponding robot (Boutilier, Goldszmidt, and Sabata 1999; Fatima 2006).² SSI auction-based coordination systems take some (but not all) synergies among targets into account and provide the following performance guarantee if the terrain and targets are initially known and do not change. Some intuition for this result can be gained from interpreting the greedy construction of minimum spanning trees as a cooperative auction (Lagoudakis et al. 2004).

Theorem 1 (Lagoudakis et al. 2005) *The sum of travel distances can be a factor of 1.5 larger than minimal but is at most a factor of two larger than minimal, whether each robot calculates its travel distance exactly or uses the cheapest insertion heuristic (Lawler et al. 1985) to determine it approximately, which results in a polynomial-time auction mechanism.*

SSI auction-based coordination systems perform even better experimentally. It can be shown that they perform hill-climbing in each round by allocating one additional target to one robot so that the sum of travel distances increases the least. This insight can be exploited to automatically determine how the robots should bid for different performance

¹Two targets have positive (negative) synergy for a robot if the smallest travel distance for visiting both targets is smaller (larger) than the sum of the smallest travel distances for visiting both targets individually from the current location of the robot.

²In practice, each robot can determine the winning bid quickly itself by listening to the bids and identifying the one with the smallest bid cost.

measures, such as having to visit all targets with a small task-completion or flow time instead of a small sum of travel distances (Tovey et al. 2005). Furthermore, every robot needs to bid on only one target per round, namely the target with the smallest bid cost, since targets with larger bid costs cannot win. The runtime of winner determination during each round is linear in the number of submitted bids, which results in the following desirable communication and runtime complexities of winner determination:

Proposition 1 (Lagoudakis et al. 2005) *The number of bids submitted by all robots during each round and the runtime of winner determination are linear in the number of robots and independent of the number of unallocated targets.*

Improving SSI Auctions

Researchers have investigated several variants of SSI auctions to build SSI auction-based coordination systems that decrease the sums of travel distances while still allocating tasks to robots in real time. Researchers have also investigated how to further improve the target allocation afterwards, for example, using task swaps among robots (Dias and Stentz 2000; Zheng and Koenig 2009).

SSI Auctions with Rollouts

Everything is the same as for SSI auction-based coordination systems, except that every robot now bids on a target the sum of travel distances of the complete target allocation that would result from it being allocated the target that it bids on in addition to all targets allocated to it in previous rounds, all other robots being allocated the targets allocated to them in previous rounds, and then hill-climbing completing this partial target allocation to a complete target allocation (Zheng, Koenig, and Tovey 2006). The bid costs of the robots are now more informed since they are based on complete rather than partial target allocations, which takes more synergies among targets into account and makes hill-climbing less myopic. This is most helpful in the early rounds where the target allocations are far from being complete. Proposition 1 continues to hold trivially. However, the number of rounds increases since each round of the main SSI auction is now preceded by the rounds of the SSI auctions for the corresponding (parallel) rollouts.

SSI Auctions with Bundle Bids

Everything is the same as for SSI auction-based coordination systems, except that every robot now bids on each bundle of at most k unallocated targets, and winner determination then allocates k additional targets to one or more robots, making SSI auctions with bundle bids the same as standard SSI auctions if $k = 1$ and the same as combinatorial auctions if k is large (Koenig et al. 2007). SSI auction-based coordination systems with bundle bids allocate k additional targets to one or more robots in each round so that the sum of travel distances increases the least, which takes more synergies among targets into account and makes hill-climbing less myopic. It can be shown that every robot needs to bid on only a constant number of bundles per round since the other

bundles cannot win. These bundles can be determined automatically. For example, every robot needs to bid on three bundles per round if $k = 2$, namely the single-target bundles with the two lowest bid costs and the double-target bundle with the lowest bid cost, and only seven bundles per round if $k = 3$, which can reduce the number of bids by several orders of magnitude. For example, the number of bids submitted by every robot per round for twenty unallocated targets and $k = 6$ is only 105 instead of 60,459 (Koenig et al. 2007). The winner determination procedure can be automatically determined and, to decrease runtime, be compiled into compact program code (Daniel and Koenig 2009). Proposition 1 continues to hold for fixed k (although the proof is nontrivial), which makes SSI auctions with bundle bids an attractive cross between SSI and combinatorial auctions (Koenig et al. 2007).

SSI Auctions with Regret Clearing

Everything is the same as for SSI auction-based coordination systems, except that winner determination now allocates the target with the largest regret to the robot whose bid cost on it is smallest, where the regret of a target is the difference of the second-smallest and smallest bid cost on it (Koenig et al. 2008). If the terrain and targets are initially known and do not change, SSI auction-based coordination systems with regret clearing provide the following performance guarantee, which could be worse than that of Theorem 1.

Theorem 2 (Koenig et al. 2008) *The sum of travel distances can be a factor of three larger than minimal but is at most a factor of twice the number of targets larger than minimal, whether each robot calculates its travel distance exactly or uses the cheapest insertion heuristic to determine it approximately, which results in a polynomial-time auction mechanism.*

However, SSI auction-based coordination systems with regret clearing can perform much better experimentally than standard SSI auction-based coordination systems, which can be explained as follows: Later allocations of targets to robots are typically more informed than earlier ones since the target allocations in earlier rounds are far from being complete. If a target is allocated to a robot in the current round then one wants to ensure that, if this allocation were postponed, the same allocation would be made in later rounds. SSI auction-based coordination systems with regret clearing achieve this objective by no longer performing hill-climbing in each round but rather maximizing the difference of the second-smallest and smallest team costs that would result from the second-best and best robot, respectively, being allocating one additional target. The communication and runtime complexities of winner determination remain polynomial but Proposition 1 no longer holds since every robot now needs to bid on all unallocated targets per round.

Extensions

Auction-based coordination systems have also been successfully applied to several generalizations of multi-robot routing problems, for example, where a team of robots has to visit a set of given targets with linear decreasing rewards

over time, such as required for the delivery of goods to rescue sites after disasters. The robots have to visit a subset of targets so as to maximize the surplus, which is defined to be the sum of the rewards of the visited targets minus the sum of travel costs. Auction-based coordination systems are able to solve these NP-hard problems in seconds and with a surplus that is comparable to the surplus found by a mixed integer program with a 12 hour time limit (Ekici, Keskinocak, and Koenig 2009). Auction-based coordination systems have also been applied to multi-robot routing problems where a team of robots has to visit a set of given targets with given priorities during given time windows that do not overlap, such as required for planetary exploration. Again, auction-based coordination systems are able to solve these NP-hard problems in seconds and with good team performance (Melvin et al. 2007). Current work includes applying SSI auction-based coordination systems to multi-robot routing problems where robots need to visit some targets simultaneously, such as required for moving large obstacles out of the way cooperatively. Every robot now needs to bid on each pair of time slot and unallocated target. The resulting function maps pairs of time slots and unallocated targets to bid costs but is typically more compact and approximated more easily if it is expressed as a function that maps unallocated targets to reaction functions, where a reaction function maps time slots to bid costs (Zheng and Koenig 2008).

Future Work

There are many alternative approaches for building coordination systems, both centralized (for example, mixed integer programming) and decentralized (for example, distributed constraint optimization or token passing). More work needs to be done on determining when to use which approach since the strengths and weaknesses of the individual approaches are not yet well understood and only a few experimental comparisons exist (Xu et al. 2006). More work needs to be done on developing auction-based coordination systems that better exploit the local (private) information of the agents and auction-based coordination systems for heterogeneous agents. More generally, more work needs to be done on applying auction-based coordination systems in more complex application domains than has been done so far.

- **Example 1:** Consider a heterogeneous team of two different kinds of agents, namely general agents X that can perform both tasks A and B, and specialized agents Y that can perform only task A. Then, it might not be a good idea to let the myopically best agent execute a task (which is what SSI auction-based coordination systems do currently). For example, if two agents X and Y are available and agent X is assigned to execute task A, then no agent is available to execute an arriving task B. On the other hand, if agent Y is assigned to execute task A, then agent X is still available to execute an arriving task A or B. Thus, the agents need to predict the future to achieve a small team cost, perhaps using methods from machine learning (Schneider et al. 2005).
- **Example 2:** Consider tasks that involve time-consuming planning or scheduling to determine the bid costs of the

agents. Then, the agents can calculate only a limited number of bid costs and first need to determine which bid costs to calculate. Thus, they need to predict the bid costs before calculating them, perhaps again using methods from machine learning (Busquets and Simmons 2006).

Finally, more work needs to be done on making auction-based coordination systems robust against error (Sariel, Balch, and Erdogan 2006; Nanjanath and Gini 2008), for example, to ensure that each target gets visited even when robots fail or leave the communication range of other robots.

References

- Berhault, M.; Huang, H.; Keskinocak, P.; Koenig, S.; and Elmaghraby, W. 2003. Robot exploration with combinatorial auctions. In *Proceedings of the International Conference on Intelligent Robots and Systems*, 1957–1962.
- Boutilier, C.; Goldszmidt, M.; and Sabata, B. 1999. Sequential auctions for the allocation of resources with complementarities. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 527–523.
- Busquets, D., and Simmons, R. 2006. Learning when to auction and when to bid. In Gini, M., and Voyles, R., eds., *Distributed Autonomous Robotic Systems 7*. Springer. 21–30.
- Clearwater, S.; Costanza, R.; Dixon, M.; and Schroeder, B. 1996. Saving energy using market-based control. In Clearwater, S., ed., *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific. 253–273.
- Daniel, K., and Koenig, S. 2009. Fast winner determination for agent coordination with SBB auctions [poster abstract]. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*, 1197–1198.
- Dias, M., and Stentz, A. 2000. A free market architecture for distributed control of a multirobot system. In *Proceedings of the International Conference on Intelligent Autonomous Systems*, 115–122.
- Dias, M.; Zlot, R.; Kalra, N.; and Stentz, A. 2006. Market-based multirobot coordination: A survey and analysis. *Proceedings of the IEEE* 94(7):1257–1270.
- Ekici, A.; Keskinocak, P.; and Koenig, S. 2009. Multi-robot routing with linear decreasing rewards over time. In *Proceedings of the International Conference on Robotics and Automation*, 958–963.
- Fatima, S. 2006. Sequential versus simultaneous auctions: A case study. In *Proceedings of the International Conference on Electronic Commerce*, 82–91.
- Frias-Martinez, V.; Sklar, E.; and Parsons, S. 2004. Exploring auction mechanisms for role assignment in teams of autonomous robots. In *Proceedings of the International RoboCup Symposium*, 532–539.
- Gerkey, B., and Mataric, M. 2002. Sold!: Auction methods for multi-robot coordination. *IEEE Transactions on Robotics and Automation* 18(5):758–768.
- Howard, A., and Viguria, A. 2007. Controlled reconfiguration of robotic mobile sensor networks using distributed

- allocation formalisms. In *Proceedings of the NASA Science Technology Conference*.
- Hunsberger, L., and Grosz, B. 2000. A combinatorial auction for collaborative planning. In *Proceedings of the International Conference on Multi-Agent Systems*, 151–158.
- Kalra, N.; Stentz, R.; and Ferguson, D. 2005. Hoplites: A market framework for planned tight coordination in multi-agent teams. In *Proceedings of the International Conference on Robotics and Automation*, 1170–1177.
- Koenig, S.; Tovey, C.; Zheng, X.; and Sungur, I. 2007. Sequential bundle-bid single-sale auction algorithms for decentralized control. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 1359–1365.
- Koenig, S.; Zheng, X.; Tovey, C.; Borie, R.; Kilby, P.; Markakis, V.; and Keskinocak, P. 2008. Agent coordination with regret clearing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 101–107.
- Lagoudakis, M.; Keskinocak, P.; Kleywegt, A.; and Koenig, S. 2004. Auctions with performance guarantees for multi-robot task allocation. In *Proceedings of the International Conference on Intelligent Robots and Systems*, 1957–1962.
- Lagoudakis, M.; Markakis, V.; Kempe, D.; Keskinocak, P.; Koenig, S.; Kleywegt, A.; Tovey, C.; Meyerson, A.; and Jain, S. 2005. Auction-based multi-robot routing. In *Proceedings of the International Conference on Robotics: Science and Systems*, 343–350.
- Lawler, E.; Lenstra, J.; Kan, A.; and Shmoys, D., eds. 1985. *The Traveling Salesman Problem*. John Wiley.
- McAfee, P., and McMillan, J. 1987. Auctions and bidding. *Journal of Economic Literature* 15:699–738.
- Melvin, J.; Keskinocak, P.; Koenig, S.; Tovey, C.; and Ozkaya, B. Y. 2007. Multi-robot routing with rewards and disjoint time windows. In *Proceedings of the International Conference on Intelligent Robots and Systems*, 958–963.
- Nair, R.; Ito, T.; Tambe, M.; and Marsella, S. 2002. Task allocation in the rescue simulation domain: A short note. In Birk, A., and Coradeschi, S., eds., *RoboCup-2001: Robot Soccer World Cup V*, Lecture Notes in Computer Science. Springer.
- Nanjanath, M., and Gini, M. 2008. Repeated auctions for robust task execution by a robot team. Technical Report 08-032, University of Minnesota, Computer Science Department, Minneapolis (Minnesota).
- Reeves, D.; Wellman, M.; MacKie-Mason, J.; and Osepashvili, A. 2005. Exploring bidding strategies for market-based scheduling. *Decision Support Systems* 39:67–85.
- Sandholm, T. 1993. An implementation of the contract net protocol based on marginal cost calculations. In *Proceedings of the International Workshop on Distributed Artificial Intelligence*, 295–308.
- Sariel, S.; Balch, T.; and Erdogan, N. 2006. Robust multi-robot cooperation through dynamic task allocation and precaution routines. In *Proceedings of the International Conference on Informatics in Control, Automation and Robotics*, 196–201.
- Sariel, S.; Balch, T.; and Stack, J. 2006. Empirical evaluation of auction-based coordination of AUVs in a realistic simulated mine countermeasure task. In *Proceedings of the International Symposium on Distributed Autonomous Robotic Systems*, 197–206.
- Schneider, J.; Apfelbaum, D.; Bagnell, D.; and Simmons, R. 2005. Learning opportunity costs in multi-robot market based planners. In *Proceedings of the International Conference on Robotics and Automation*, 1151–1156.
- Simmons, R.; Apfelbaum, D.; Burgard, W.; Fox, D.; Moors, M.; Thrun, S.; and Younes, H. 2000. Coordination for multi-robot exploration and mapping. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 852–858.
- Smith, R. 1980. The contract net protocol: high level communication and control in a distributed problem solver. *IEEE Transactions on Computers* C-29:1104–1113.
- Sycara, K.; Scerri, P.; Srinivas, S.; and Lewis, M. 2005. Task allocation in teams for launch and range operations. In *Proceedings of the International Conference on Human Computer Interaction*.
- Thayer, S.; Digney, B.; Dias, M.; Stentz, A.; Nabbe, B.; and Hebert, M. 2000. Distributed robotic mapping of extreme environments. In *Proceedings of SPIE: Mobile Robots XV and Telemanipulator and Telepresence Technologies VII*, volume 4195, 84–95.
- Tovey, C.; Lagoudakis, M.; Jain, S.; and Koenig, S. 2005. The generation of bidding rules for auction-based robot coordination. In Parker, L.; Schneider, F.; and Schultz, A., eds., *Multi-Robot Systems: From Swarms to Intelligent Automata*. Springer. 3–14.
- Xu, Y.; Scerri, P.; Sycara, K.; and Lewis, M. 2006. Comparing market and token-based coordination [short paper]. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*, 1113–1115.
- Yu, Y., and Prasanna, V. 2005. Energy balanced task allocation for collaborative processing in wireless sensor networks. *Mobile Networks and Applications* 10(1-2):115–131.
- Zheng, X., and Koenig, S. 2008. Reaction functions for task allocation to cooperative agents. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*, 559–566.
- Zheng, X., and Koenig, S. 2009. K-swaps: Cooperative negotiation for solving task-allocation problems. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 373–379.
- Zheng, X.; Koenig, S.; and Tovey, C. 2006. Improving sequential single-item auctions. In *Proceedings of the International Conference on Intelligent Robots and Systems*, 2238–2244.
- Zlot, R.; Stentz, A.; Dias, M.; and Thayer, S. 2002. Multi-robot exploration controlled by a market economy. In *Proceedings of the International Conference on Robotics and Automation*, 3016–3023.